Los Alamos, Hewlett Packard Enterprise and NVIDIA partner to speed up scientific computing

October 6, 2020

LOS ALAMOS, N.M., Oct. 6, 2020—Los Alamos National Laboratory is partnering with Hewlett Packard Enterprise (HPE) and NVIDIA to focus on delivering next-generation technologies to accelerate scientific computing. New developments will include innovative high-performance computing (HPC) technology efforts to advance greater performance efficiency, workflow efficiency and analytics. Additional details on the collaboration will be unveiled this fall.

"We are excited to be part of this long-term collaboration. Together, we can accelerate the adoption of new HPC advancements and technologies that enable efficient computing for future HPC applications," said John Sarrao, Deputy Laboratory Director for Science, Technology and Engineering at Los Alamos.

All three organizations bring rich histories of technology innovation to the partnership, building upon more than a decade of deep collaborations with the Department of Energy that have pushed the boundaries of what is possible in HPC.

Bob Webster, Deputy Laboratory Director for Weapons at Los Alamos, expanded on the importance of these collaborations. "A big challenge for many HPC users is the effort to port their workloads to run on these new exascale and beyond supercomputers. This long-term partnership, with its interest in next-generation codesign, will raise the portability efficiency and benefit the entire HPC user community, not just those with high-fidelity 3D simulation demands such as Los Alamos.' "

The Laboratory has made significant contributions through its world-class leadership in HPC architectures and storage including Roadrunner and its predecessor design project Darkhorse—all of which foresaw the common addressability between CPUs and accelerators in 2001 and produced the first petaflop heterogeneous HPC cluster in the world in 2008. Additionally, Los Alamos' burst buffer and campaign storage innovations have netted huge payoffs for more efficient data movement for all high-end HPC users.

Similarly, HPE has a long history in driving innovation and advancing discovery to solve the world's toughest problems with HPC by delivering a powerful portfolio of HPC and AI solutions. HPE's offerings, which integrate technologies from Cray and span across compute, networking, software, storage and services, have been architected from the ground-up to support the exascale era and the growing need for emerging AI, machine learning and analytics applications with unprecedented performance and efficiency.

"We are committed to advancing a diverse set of flexible solutions and options, beyond the innovative technologies employed in the initial exaclass machines, for our growing, broad range of research and engineering needs," said Bill Mannel, vice president and general manager of HPC at HPE. "Our new long-term collaboration with Los Alamos National Laboratory and NVIDIA will advance next-generation HPC architectures and designs, and significantly unlock new value in improved workload performance, increased compute efficiency, performance, and analytical capabilities."

NVIDIA's invention of the GPU in 1999 revolutionized parallel computing, and more recently, ignited modern AI. Today, NVIDIA's GPU-accelerated computing platform powers the fastest supercomputers in the U.S. and Europe, and several of the world's most advanced systems under development, enabling the performance and computing efficiency required to solve the world's most pressing challenges, including addressing climate change, discovering new drugs and uncovering new energy sources. NVIDIA also brings its comprehensive development ecosystem that includes deep expertise in scientific computing, molecular dynamics, quantum chemistry and multi-physics in more than 1,000 HPC applications.

"The greatest obstacle in advancing supercomputing today is scaling energy efficiency with the growing diversity of application workloads and performance demands required by modern HPC challenges," said Ian Buck, vice president and general manager of Accelerated Computing at NVIDIA. "Finding opportunities for the organizations defining the future of supercomputing to work together is how we will most efficiently move supercomputing forward."